arXiv:1608.03680v1 [cs.CG] 12 Aug 2016

The (1|1)-Centroid Problem on the Plane Concerning Distance Constraints *

Hung-I Yu, Tien-Ching Lin, and D. T. Lee

Institute of Information Science, Academia Sinica,
Nankang, Taipei 115, Taiwan,
{herbert,kero,dtlee}@iis.sinica.edu.tw

Abstract. In 1982, Drezner proposed the (1|1)-centroid problem on the plane, in which two players, called the leader and the follower, open facilities to provide service to customers in a competitive manner. The leader opens the first facility, and the follower opens the second. Each customer will patronize the facility closest to him (ties broken in favor of the first one), thereby decides the market share of the two facilities. The goal is to find the best position for the leader's facility so that its market share is maximized. The best algorithm of this problem is an $O(n^2 \log n)$-time parametric search approach, which searches over the space of market share values.

In the same paper, Drezner also proposed a general version of (1|1)-centroid problem by introducing a minimal distance constraint R, such that the follower's facility is not allowed to be located within a distance R from the leader's. He proposed an $O(n^5 \log n)$-time algorithm for this general version by identifying $O(n^4)$ points as the candidates of the optimal solution and checking the market share for each of them. In this paper, we develop a new parametric search approach searching over the $O(n^4)$ candidate points, and present an $O(n^2 \log n)$-time algorithm for the general version, thereby close the $O(n^3)$ gap between the two bounds.

Keywords: competitive facility, Euclidean plane, parametric search

1 Introduction

In 1929, economist Hotelling introduced the first competitive location problem in his seminal paper [13]. Since then, the subject of competitive facility location has been extensively studied by researchers in the fields of spatial economics, social and political sciences, and operations research, and spawned hundreds of contributions in the literature. The interested reader is referred to the following survey papers [3,7,8,9,11,12,17,19].

Hakimi [10] and Drezner [5] individually proposed a series of competitive location problems in a leader-follower framework. The framework is briefly described as follows. There are n customers in the market, and each is endowed

* Research supported by under Grants No. MOST 103-2221-E-005-042, 103-2221-E-005-043.

with a certain buying power. Two players, called the *leader* and the *follower*, sequentially open facilities to attract the buying power of customers. At first, the leader opens his p facilities, and then the follower opens another r facilities. Each customer will patronize the closest facility with all buying power (ties broken in favor of the leader's ones), thereby decides the market share of the two players. Since both players ask for market share maximization, two competitive facility location problems are defined under this framework. Given that the leader locates his p facilities at the set X_p of p points, the follower wants to locate his r facilities in order to attract the most buying power, which is called the $(r|X_p)$-*medianoid* problem. On the other hand, knowing that the follower will react with maximization strategy, the leader wants to locate his p facilities in order to retain the most buying power against the competition, which is called the $(r|p)$-*centroid* problem.

Drezner [5] first proposed to study the two competitive facility location problems on the Euclidean plane. Since then, many related results [4,5,6,11,14] have been obtained for different values of r and p. Due to page limit, here we introduce only previous results about the case $r = p = 1$. For the $(1|X_1)$-medianoid problem, Drezner [5] showed that there exists an optimal solution arbitrarily close to X_1, and solved the problem in $O(n \log n)$ time by sweeping technique. Later, Lee and Wu [14] obtained an $\Omega(n \log n)$ lower bound for the $(1|X_1)$-medianoid problem, and thus proved the optimality of the above result. For the $(1|1)$-centroid problem, Drezner [5] developed a parametric search based approach that searches over the space of $O(n^2)$ possible market share values, along with an $O(n^4)$-time test procedure constructing and solving a linear program of $O(n^2)$ constraints, thereby gave an $O(n^4 \log n)$-time algorithm. Then, by improving the test procedure via Megiddo's result [16] for solving linear programs, Hakimi [11] reduced the time complexity to $O(n^2 \log n)$.

In [5], Drezner also proposed a more general setting for the leader-follower framework by introducing a *minimal distance constraint* $R \geq 0$ into the $(1|X_1)$-medianoid problem and the $(1|1)$-centroid problem, such that the follower's facility is not allowed to be located within a distance R from the leader's. The augmented problems are respectively called the $(1|X_1)_R$-*medianoid* problem and $(1|1)_R$-*centroid* problem in this paper. Drezner showed that the $(1|X_1)_R$-medianoid problem can also be solved in $O(n \log n)$ time by using nearly the same proof and technique as for the $(1|X_1)$-medianoid problem. However, for the $(1|1)_R$-centroid problem, he argued that it is hard to generalize the approach for the $(1|1)$-centroid problem to solve this general version, due to the change of problem properties. Then, he gave an $O(n^5 \log n)$-time algorithm by identifying $O(n^4)$ candidate points on the plane, which contain at least one optimal solution, and performing medianoid computation on each of them. So far, the $O(n^3)$ bound gap between the two centroid problems remains unclosed.

In this paper, we propose an $O(n^2 \log n)$-time algorithm for the $(1|1)_R$-centroid problem on the Euclidean plane, thereby close the gap last for decades. Instead of searching over market share values, we develop a new approach based on the parametric search technique by searching over the $O(n^4)$ candidate points

mentioned in [5]. This is made possible by making a critical observation on the distribution of optimal solutions for the $(1|X_1)_R$-medianoid problem given X_1, which provides us a useful tool to prune candidate points with respect to X_1. We then extend the usage of this tool to design a key procedure to prune candidates with respect to a given vertical line.

The rest of this paper is organized as follows. Section 2 gives formal problem definitions and describes previous results in [5,11]. In Section 3, we make the observation on the $(1|X_1)_R$-medianoid problem, and make use of it to find a "local" centroid on a given line. This result is then extended as a new pruning procedure with respect to any given line in Section 4, and utilized in our parametric search approach for the $(1|1)_R$-centroid problem. Finally, in Section 5, we give some concluding remarks.

2 Notations and Preliminary Results

Let $V = \{v_1, v_2, \cdots, v_n\}$ be a set of n points on the Euclidean plane $\mathbb{R}^2$, as the representatives of the n customers. Each point $v_i \in V$ is assigned with a positive weight $w(v_i)$, representing its buying power. To simplify the algorithm description, we assume that the points in V are in general position, that is, no three points are collinear and no two points share a common x or y-coordinate.

Let $d(u, w)$ denote the Euclidean distance between any two points $u, w \in \mathbb{R}^2$. For any set Z of points on the plane, we define $W(Z) = \sum\{w(v)|v \in V \bigcap Z\}$. Suppose that the leader has located his facility at $X_1 = \{x\}$, which is shortened as x for simplicity. Due to the minimal distance constraint R mentioned in [5], any point $y' \in \mathbb{R}^2$ with $d(y', x) < R$ is infeasible to be the follower's choice. If the follower locates his facility at some feasible point y, the set of customers patronizing y instead of x is defined as $V(y|x) = \{v \in V|d(v, y) < d(v, x)\}$, with their total buying power $W(y|x) = W(V(y|x))$. Then, the largest market share that the follower can capture is denoted by the function

$$W^*(x) = \max_{y \in \mathbb{R}^2, d(y,x) \geq R} W(y|x),$$

which is called the *weight loss* of x. Given a point $x \in \mathbb{R}^2$, the $(1|x)_R$-medianoid problem is to find a $(1|x)_R$-*medianoid*, which denotes a feasible point $y^* \in \mathbb{R}^2$ maximizing the weight loss of x.

In contrast, the leader tries to minimize the weight loss of his own facility by finding a point $x^* \in \mathbb{R}^2$ such that

$$W^*(x^*) \leq W^*(x)$$

for any point $x \in \mathbb{R}^2$. The $(1|1)_R$-centroid problem is to find a $(1|1)_R$-*centroid*, which denotes a point x^* minimizing its weight loss. Note that, when $R = 0$, the two problems degenerate to the $(1|x)$-medianoid and $(1|1)$-centroid problems.

2.1 Previous approaches

In this subsection, we briefly review previous results for the $(1|x)_R$-medianoid, $(1|1)$-centroid, and $(1|1)_R$-centroid problems in [5,11], so as to derive some basic properties essential to our approach.

Let L be an arbitrary line, which partitions the Euclidean plane into two half-planes. For any point $y \notin L$, we define $H(L, y)$ as the close half-plane including y, and $H^-(L, y)$ as the open half-plane including y (but not L). For any two distinct points $x, y \in \mathbb{R}^2$, let $B(y|x)$ denote the perpendicular bisector of the line segment from x to y.

Given an arbitrary point $x \in \mathbb{R}^2$, we first describe the algorithm for finding a $(1|x)_R$-medianoid in [5]. Let y be an arbitrary point other than x, and y' be some point on the open line segment from y to x. We can see that $H^-(B(y|x), y) \subset H^-(B(y'|x), y')$, which implies the fact that $W(y'|x) = W(H^-(B(y'|x), y')) \geq W(H^-(B(y|x), y)) = W(y|x)$, It shows that moving y toward x does not diminish its weight capture, thereby follows the lemma.

Lemma 1 *[5] There exists a $(1|x)_R$-medianoid in $\{y \mid y \in \mathbb{R}^2, d(x, y) = R\}$.*

For any point $z \in \mathbb{R}^2$, let $C_R(x)$ and $C_\gamma(x)$ be the circles centered at z with radii R and $\gamma = R/2$, respectively. By Lemma 1, finding a $(1|x)_R$-medianoid can be reduced to searching a point y on $C_R(x)$ maximizing $W(y|x)$. Since the perpendicular bisector $B(y|x)$ of each point y on $C_R(x)$ is a tangent line to the circle $C_\gamma(x)$, the searching of y on $C_R(x)$ is equivalent to finding a tangent line to $C_\gamma(x)$ that partitions the most weight from x. The latter problem can be solved in $O(n \log n)$ time as follows. For each $v \in V$ outside $C_\gamma(x)$, we calculate its two tangent lines to $C_\gamma(x)$. Then, by sorting these tangent lines according to the polar angles of their corresponding tangent points with respect to x, we can use the angle sweeping technique to check how much weight they partition.

Theorem 1 *[5] Given a point $x \in \mathbb{R}^2$, the $(1|x)_R$-medianoid problem can be solved in $O(n \log n)$ time.*

Next, we describe the algorithm of the $(1|1)_R$-centroid problem in [5]. Let S be a subset of V. We define $\mathcal{C}(S)$ to be the set of all circles $C_\gamma(v)$, $v \in S$, and $CH(\mathcal{C}(S))$ to be the convex hull of these circles. It is easy to see the following.

Lemma 2 *[5] Let S be a subset of V. For any point $x \in \mathbb{R}^2$, $W^*(x) \geq W(S)$ if x is outside $CH(\mathcal{C}(S))$.*

For any positive number W_0, let $I(W_0)$ be the intersection of all convex hulls $CH(\mathcal{C}(S))$, where $S \subseteq V$ and $W(S) \geq W_0$. We have the lemma below.

Lemma 3 *[5] Let W_0 be a positive real number. For any point $x \in \mathbb{R}^2$, $W^*(x) < W_0$ if and only if $x \in I(W_0)$.*

Proof. Consider first the case that $x \in I(W_0)$. By definition, x intersects with every $CH(\mathcal{C}(S))$ of subset $S \subseteq V$ with $W(S) \geq W_0$. Let $S' \subseteq V$ be any of such subsets. Since $x \in CH(\mathcal{C}(S'))$, for any point y feasible to x, there must exist a point $v \in S$ such that $v \notin H^-(B(y|x), y)$, implying that no feasible point y can acquire all buying power from customers of S'. It follows that no feasible point y can acquire buying power larger than or equal to W_0, i.e., $W^*(x) < W_0$.

If $x \notin I(W_0)$, there must exist a subset $S \subseteq V$ with $W(S) \geq W_0$, such that $x \notin CH(\mathcal{C}(S))$. By Lemma 2, $W^*(x) \geq W(S) \geq W_0$. □

Drezner [5] argued that the set of all $(1|1)_R$-centroids is equivalent to some intersection $I(W_0)$ for smallest possible W_0. We slightly strengthen his argument below. Let $\mathcal{W} = \{W(y|x) \mid x, y \in \mathbb{R}^2, d(x, y) \geq R\}$. The following lemma can be obtained.

Lemma 4. *Let W_0^* be the smallest number in $\mathcal{W}$ such that $I(W_0^*)$ is not null. A point x is a $(1|1)_R$-centroid if and only if $x \in I(W_0^*)$.*

Proof. Let W_{OPT} be the weight loss of some $(1|1)_R$-centroid x^*. We first show that $I(W_0)$ is null for any $W_0 \leq W_{OPT}$. Suppose to the contrary that it is not null and there exists a point x' in $I(W_0)$. By Lemma 3, $W^*(x') < W_0 \leq W_{OPT}$, which contradicts the optimality of x^*. Moreover, since $I(W_0^*)$ is not null, we have that $W_0^* > W_{OPT}$.

We now show that a point x is a $(1|1)_R$-centroid if and only if $x \in I(W_0^*)$. If x is a $(1|1)_R$-centroid, we have that $W^*(x) = W_{OPT} < W_0^*$. By Lemma 3, $x \in I(W_0^*)$. On the other hand, if x is not a $(1|1)_R$-centroid, we have that $W^*(x) > W_{OPT}$. Since by definition $W^*(x) \in \mathcal{W}$, we can see that $W^*(x) \geq W_0^*$. Thus, again by Lemma 3, $x \notin I(W_0^*)$. □

Although it is hard to compute $I(W_0^*)$ itself, we can find its vertices as solutions to the $(1|1)_R$-centroid problem. Let $\mathcal{T}$ be the set of outer tangent lines of all pairs of circles in $\mathcal{C}(V)$. For any subset $S \subseteq V$, the boundary of $CH(\mathcal{C}(S))$ is formed by segments of lines in $\mathcal{T}$ and arcs of circles in $\mathcal{C}(V)$. Since $I(W_0)$ is an intersection of such convex hulls, its vertices must fall within the set of intersection points between lines in $\mathcal{T}$, between circles in $\mathcal{C}(V)$, and between one line in $\mathcal{T}$ and one circle in $\mathcal{C}(V)$. Let $\mathcal{T} \times \mathcal{T}$, $\mathcal{C}(V) \times \mathcal{C}(V)$, and $\mathcal{T} \times \mathcal{C}(V)$ denote the three sets of intersection points, respectively. We have the lemma below.

Lemma 5 *[5] There exists a $(1|1)_R$-centroid in $\mathcal{T} \times \mathcal{T}$, $\mathcal{C}(V) \times \mathcal{C}(V)$, and $\mathcal{T} \times \mathcal{C}(V)$.*

Obviously, there are at most $O(n^4)$ intersection points, which can be viewed as the *candidates* of being $(1|1)_R$-centroids. Drezner thus gave an algorithm by evaluating the weight loss of each candidate by Theorem 1.

Theorem 2 *[5] The $(1|1)_R$-centroid problem can be solved in $O(n^5 \log n)$ time.*

We remark that, when $R = 0$, $CH(\mathcal{C}(S))$ for any $S \subseteq V$ degenerates to a convex polygon, so does $I(W_0)$ for any given W_0, if not null. Drezner [5] proved

that in this case $I(W_0)$ is equivalent to the intersection of all half-planes H with $W(H) \geq W_0$. Thus, whether $I(W_0)$ is null can be determined by constructing and solving a linear program of $O(n^2)$ constraints, which takes $O(n^2)$ time by Megiddo's result [16]. Since $|\mathcal{W}| = O(n^2)$, by Lemma 4, the $(1|1)$-centroid problem can be solved in $O(n^2 \log n)$ time [11], by applying parametric search over $\mathcal{W}$ for W_0^*. Unfortunately, it is hard to generalize this idea to the case $R > 0$, motivating us to develop a different approach.

3 Local $(1|1)_R$-Centroid within a Line

In this section, we analyze the properties of $(1|x)_R$-medianoids of a given point x in Subsection 3.1, and derive a procedure that prunes candidate points with respect to x. Applying this procedure, we study a restricted version of the $(1|1)_R$-centroid problem in Subsection 3.2, in which the leader's choice is limited to a given non-horizontal line L, and obtain an $O(n \log^2 n)$-time algorithm. The algorithm is then extended as the basis of the test procedure for the parametric search approach in Section 4.

3.1 Pruning with Respect to a Point

Given a point $x \in \mathbb{R}^2$ and an angle θ between 0 and 2π, let $y(\theta|x)$ be the point on $C_R(x)$ with polar angle θ with respect to x.[1] We define $MA(x) = \{\theta \mid W(y(\theta|x)|x) = W^*(x), 0 \leq \theta < 2\pi\}$, that is, the set of angles θ maximizing $W(y(\theta|x)|x)$ (see Figure 1). It can be observed that, for any $\theta \in MA(x)$ and sufficiently small ϵ, both $\theta + \epsilon$ and $\theta - \epsilon$ belong to $MA(x)$, because each $v \in V(y(\theta|x)|x)$ does not intersect $B(y(\theta|x)|x)$ by definition. This implies that angles in $MA(x)$ form open angle interval(s) of non-zero length.

To simplify the terms, let $W(\theta|x) = W(y(\theta|x)|x)$ and $B(\theta|x) = B(y(\theta|x)|x)$ in the remaining of this section. Also, let $F(\theta|x)$ be the line passing through x and parallel to $B(\theta|x)$. The following lemma provides the basis for pruning.

Lemma 6 *Let $x \in \mathbb{R}^2$ be an arbitrary point, and θ be an angle in $MA(x)$. For any point $x' \notin H^-(F(\theta|x), y(\theta|x))$, $W^*(x') \geq W^*(x)$.*

Proof. Since $x' \notin H^-(F(\theta|x), y(\theta|x))$ and $y(\theta|x) \in H^-(F(\theta|x), y(\theta|x))$, by the definition of bisectors, the distance between $F(\theta|x')$ and $B(\theta|x)$ is no less than $R/2$, which implies that $H^-(B(\theta|x), y(\theta|x)) \subseteq H^-(B(\theta|x'), y(\theta|x'))$. Therefore, we can derive the following inequality

$$\begin{aligned} W^*(x') &\geq W(\theta|x') \\ &= W(H^-(B(\theta|x'), y(\theta|x')) \\ &\geq W(H^-(B(\theta|x), y(\theta|x)) \\ &= W(\theta|x) \\ &= W^*(x), \end{aligned}$$

[1] We assume that a polar angle is measured counterclockwise from the positive x-axis.

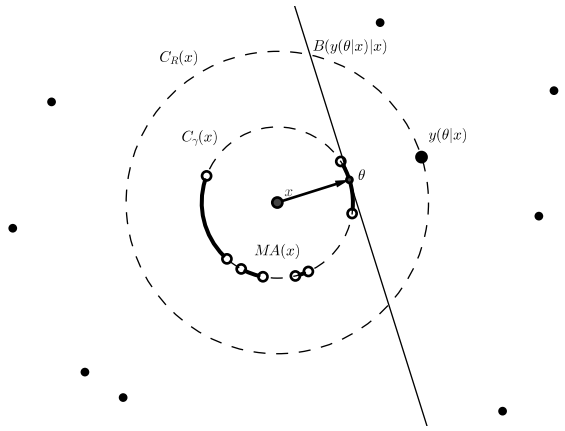

Fig. 1. The black arcs represent the intervals of angles in $MA(x)$, whereas the open circles represent the open ends of these intervals.

which completes the proof. □

This lemma tells us that, given a point x and an angle $\theta \in MA(x)$, all points not in $H^-(F(\theta|x), y(\theta|x))$ can be ignored while finding $(1|1)_R$-centroids, as their weight losses are no less than that of x. By this lemma, we can also prove that the weight loss function is convex along any line on the plane, as shown below.

Lemma 7 *Let x_1, x_2 be two arbitrary distinct points on a given line L. For any point $x \in \overline{x_1x_2}\backslash\{x_1, x_2\}$, $W^*(x) \leq \max\{W^*(x_1), W^*(x_2)\}$.*

Proof. Suppose by contradiction that $W^*(x) > W^*(x_1)$ and $W^*(x) > W^*(x_2)$ for some point $x \in \overline{x_1x_2}\backslash\{x_1, x_2\}$. Since $W^*(x) > W^*(x_1)$, by Lemma 6 there exists an angle $\theta \in MA(x)$ such that x_1 is included in $H^-(F(\theta|x), y(\theta|x))$. However, since $x \in \overline{x_1x_2}\backslash\{x_1, x_2\}$, x_1 and x_2 locate on different sides of $F(\theta|x)$. It follows that x_2 is outside $H^-(F(\theta|x), y(\theta|x))$ and $W^*(x_2) \geq W^*(x)$ by Lemma 6, which contradicts the assumption. Thus, the lemma holds. □

We further investigate the distribution of angles in $MA(x)$. Let $CA(x)$ be the minimal angle interval covering all angles in $MA(x)$ (see Figure 2(a)), and $\delta(CA(x))$ be its angle span in radians. As mentioned before, $MA(x)$ consists of open angle interval(s) of non-zero length, which implies that $CA(x)$ is an open interval and $\delta(CA(x)) > 0$. Moreover, we can derive the following.

Lemma 8 *If $\delta(CA(x)) > \pi$, x is a $(1|1)_R$-centroid.*

Proof. We prove this lemma by showing that $W^*(x') \geq W^*(x)\ \forall\ x' \neq x$. Let $x' \in \mathbb{R}^2$ be an arbitrary point other than x, and θ' be its polar angle with respect to x. Obviously, any angle θ satisfying $x' \in H^-(F(\theta|x), y(\theta|x))$ is in the open interval $(\theta' - \pi/2, \theta' + \pi/2)$, the angle span of which is equal to π. Since

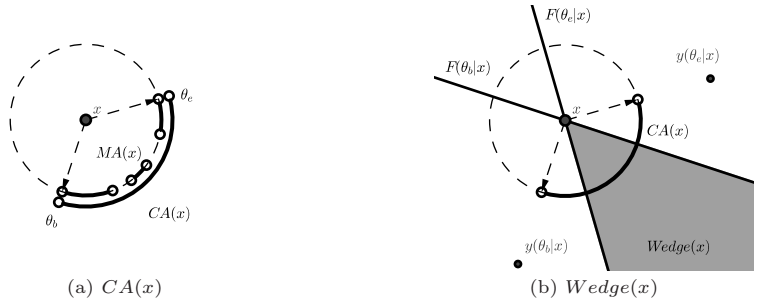

(a) $CA(x)$ (b) $Wedge(x)$

Fig. 2. $CA(x)$ and $Wedge(x)$.

$\delta(CA(x)) > \pi$, by its definition there exists an angle $\theta \in MA(x)$ such that $x' \notin H^-(F(\theta|x), y(\theta|x))$. Thus, by Lemma 6, we have $W^*(x') \geq W^*(x)$, thereby proves the lemma. □

We call a point x satisfying Lemma 8 a *strong* $(1|1)_R$*-centroid*, since its discovery gives an immediate solution to the $(1|1)_R$-centroid problem. Note that there are problem instances in which no strong $(1|1)_R$-centroids exist.

Suppose that $\delta(CA(x)) \leq \pi$ for some point $x \in \mathbb{R}^2$. Let $Wedge(x)$ denote the *wedge* of x, defined as the intersection of the two half-planes $H(F(\theta_b|x), y(\theta_b|x))$ and $H(F(\theta_e|x), y(\theta_e|x))$, where θ_b and θ_e are the beginning and ending angles of $CA(x)$, respectively. As illustrated in Figure 2(b), $Wedge(x)$ is the infinite region lying between two half-lines extending from x (including x and the two half-lines). The half-lines defined by $F(\theta_e|x)$ and $F(\theta_b|x)$ are called its boundaries, and the counterclockwise (CCW) angle between the two boundaries is denoted by $\delta(Wedge(x))$. Since $0 < \delta(CA(x)) \leq \pi$, we have that $Wedge(x) \neq \emptyset$ and $0 \leq \delta(Wedge(x)) < \pi$.

It should be emphasized that $Wedge(x)$ is a computational byproduct of $CA(x)$ when x is not a strong $(1|1)_R$-centroid. In other words, not every point has its wedge. Therefore, we make the following assumption (or restriction) in order to avoid the misuse of $Wedge(x)$.

Assumption 1 *Whenever* $Wedge(x)$ *is mentioned, the point* x *has been found not to be a strong* $(1|1)_R$*-centroid, either by computation or by properties. Equivalently,* $\delta(CA(x)) \leq \pi$.

The following essential lemma makes $Wedge(x)$ our main tool for prune-and-search. (Note that its proof cannot be trivially derived from Lemma 6, since by definition θ_b and θ_e do not belong to the open intervals $CA(x)$ and $MA(x)$.)

Lemma 9 *Let* $x \in \mathbb{R}^2$ *be an arbitrary point. For any point* $x' \notin Wedge(x)$, $W^*(x') \geq W^*(x)$.

Proof. By symmetry, suppose that $x' \notin H(F(\theta_b|x), y(\theta_b|x))$. We can further divide the position of x' into two cases, (1) $x' \in H(F(\theta_e|x), y(\theta_e|x))$ and (2) $x' \notin H(F(\theta_e|x), y(\theta_e|x))$.

Consider case (1). The two assumptions ensure that there exists an angle $\theta' \in (\theta_b, \theta_e]$, such that $F(\theta'|x)$ passes through x'. Obviously, any angle $\theta'' \in (\theta_b, \theta')$ satisfies that $x' \notin H(F(\theta''|x), y(\theta''|x))$. By the definition of $CA(x)$, there must exist an angle $\theta'_b \in (\theta_b, \theta')$ infinitely close to θ_b, such that θ'_b belongs to $MA(x)$. Thus, by Lemma 6, we have that $W^*(x') \geq W^*(x)$.

In case (2), for any angle $\theta'' \in MA(x)$, we have that $x' \notin H(F(\theta''|x), y(\theta''|x))$, since θ'' is in (θ_b, θ_e). Again, $W^*(x') \geq W^*(x)$ by Lemma 6. □

Finally, we consider the computation of $Wedge(x)$.

Lemma 10 *Given a point $x \in \mathbb{R}^2$, $MA(x)$, $CA(x)$, and $Wedge(x)$ can be computed in $O(n \log n)$ time.*

Proof. By Theorem 1, we first compute $W^*(x)$ and those ordered tangent lines in $O(n \log n)$ time. Then, by performing angle sweeping around $C_\gamma(x)$, we can identify in $O(n)$ time those open intervals of angles θ with $W(\theta|x) = W^*(x)$, of which $MA(x)$ consists. Again by sweeping around $C_\gamma(x)$, $CA(x)$ can be obtained from $MA(x)$ in $O(n)$ time. Now, if we find x to be a strong $(1|1)_R$-centroid by checking $\delta(CA(x))$, the $(1|1)_R$-centroid problem is solved and the algorithm can be terminated. Otherwise, $Wedge(x)$ can be constructed in $O(1)$ time. □

3.2 Searching on a Line

Although computing wedges can be used to prune candidate points, it does not serve as a stable prune-and-search tool, since wedges of different points have indefinite angle intervals and spans. However, Assumption 1 makes it work fine with lines. Here we show how to use the wedges to compute a *local optimal* point on a given line, i.e. a point x with $W^*(x) \leq W^*(x')$ for any point x' on the line.

Let L be an arbitrary line, which is assumed to be non-horizontal for ease of discussion. For any point x on L, we can compute $Wedge(x)$ and make use of it for pruning purposes by defining its *direction* with respect to L. Since $\delta(Wedge(x)) < \pi$ by definition, there are only three categories of directions according to the intersection of $Wedge(x)$ and L:

Upward – the intersection is the half-line of L above and including x;
Downward – the intersection is the half-line of L below and including x;
Sideward – the intersection is x itself.

If $Wedge(x)$ is sideward, x is a local optimal point on L, since by Lemma 9 $W^*(x) \leq W^*(x') \; \forall \; x' \in L$. Otherwise, either $Wedge(x)$ is upward or downward, the points on the opposite half of L can be pruned by Lemma 9. It shows that computing wedges acts as a predictable tool for pruning on L.

Next, we list sets of *breakpoints* on L in which a local optimal point locates. Recall that $\mathcal{T}$ is the set of outer tangent lines of all pairs of circles in $\mathcal{C}(V)$. We define the *$\mathcal{T}$-breakpoints* as the set $L \times \mathcal{T}$ of intersection points between L and lines in $\mathcal{T}$, and the *$\mathcal{C}$-breakpoints* as the set $L \times \mathcal{C}(V)$ of intersection points between L and circles in $\mathcal{C}(V)$. We have the following lemmas for breakpoints.

Lemma 11 *Let x_1, x_2 be two distinct points on L. If $W^*(x_1) > W^*(x_2)$, there exists at least a breakpoint on the segment $\overline{x_1x_2}\backslash\{x_1\}$.*

Proof. Let θ be an arbitrary angle in $MA(x_1)$ and S be the subset of V located in the half-plane $H^-(B(\theta|x_1), y(\theta|x_1))$. By definition, x_1 is outside the convex hull $CH(\mathcal{C}(S))$ and $W^*(x_1) = W(S)$. On the other hand, since $W^*(x_2) < W^*(x_1) = W(S)$ by assumption, we have that x_2 is inside $CH(\mathcal{C}(S))$ by Lemma 2. Thus, the segment $\overline{x_1x_2}\backslash\{x_1\}$ intersects with the boundary of $CH(\mathcal{C}(S))$. Since the boundary of $CH(\mathcal{C}(S))$ consists of segments of lines in $\mathcal{T}$ and arcs of circles in $\mathcal{C}(V)$, the intersection point is either a $\mathcal{T}$-breakpoint or a $\mathcal{C}$-breakpoint, thereby proves the lemma. □

Lemma 12 *There exists a local optimal point x_L^* which is also a breakpoint.*

Proof. Let x_L^* be a local optimal point such that $W^*(x') > W^*(x_L^*)$ for some point x' adjacent to x_L^* on L. Note that, if no such local optimal point exists, every point on L must have the same weight loss and be local optimal, and the lemma holds trivially. If such x_L^* and x' exist, by Lemma 11 there is a breakpoint on $\overline{x'x_L^*}\backslash\{x'\}$, which is x_L^* itself. Thus, the lemma holds. □

We remark that outer tangent lines parallel to L are exceptional cases while considering breakpoints. For any line $T \in \mathcal{T}$ that is parallel to L, either T does not intersect with L or they just coincide. In either case, T is irrelevant to the finding of local optimal points, and should not be counted for defining $\mathcal{T}$-breakpoints.

Now, by Lemma 12, if we have all breakpoints on L sorted in the decreasing order of their y-coordinates, a local optimal point can be found by performing binary search using wedges. Obviously, such sorted sequence can be obtained in $O(n^2 \log n)$ time, since $|L \times \mathcal{T}| = O(n^2)$ and $|L \times \mathcal{C}(V)| = O(n)$. However, in order to speed up the computations of local optimal points on multiple lines, alternatively we propose an $O(n^2 \log n)$-time preprocessing, so that a local optimal point on any given line can be computed in $O(n \log^2 n)$ time.

The preprocessing itself is very simple. For each point $v \in V$, we compute a sequence $P(v)$, consisting of points in $V\backslash\{v\}$ sorted in increasing order of their polar angles with respect to v. The computation for all $v \in V$ takes $O(n^2 \log n)$ time in total. Besides, all outer tangent lines in $\mathcal{T}$ are computed in $O(n^2)$ time. We will show that, for any given line L, $O(n)$ sorted sequences can be obtained from these pre-computed sequences in $O(n \log n)$ time, which can be used to replace a sorted sequence of all $\mathcal{T}$-breakpoints in the process of binary search.

For any two points $v \in V$ and $z \in \mathbb{R}^2$, let $T^r(z|v)$ be the outer tangent line of $C_\gamma(v)$ and $C_\gamma(z)$ to the right of the line from v to z. Similarly, let $T^l(z|v)$ be the outer tangent line to the left. (See Figure 3.) Moreover, let $t_L^r(z|v)$ and $t_L^l(z|v)$ be the points at which $T^r(z|v)$ and $T^l(z|v)$ intersect with L, respectively. We partition $\mathcal{T}$ into $O(n)$ sets $\mathcal{T}^r(v) = \{T^r(v_i|v)|v_i \in V\backslash\{v\}\}$ and $\mathcal{T}^l(v) = \{T^l(v_i|v)|v_i \in V\backslash\{v\}\}$ for $v \in V$, and consider their corresponding $\mathcal{T}$-breakpoints independently. By symmetry, we only discuss the case about $L \times \mathcal{T}^r(v)$.

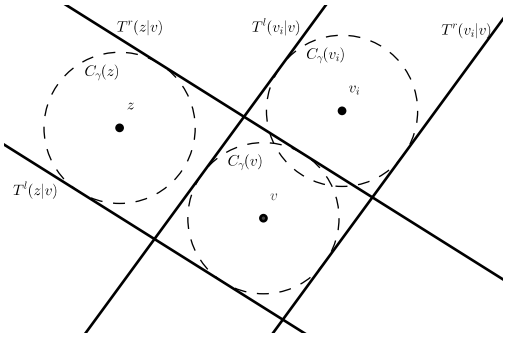

Fig. 3. Outer tangent lines of v.

Lemma 13 *For each $v \in V$, we can compute $O(1)$ sequences of $\mathcal{T}$-breakpoints on L, which satisfy the following conditions:*

(a) Each sequence is of length $O(n)$ and can be obtained in $O(\log n)$ time.
(b) Breakpoints in each sequence are sorted in decreasing y-coordinates.
(c) The union of breakpoints in all sequences form $L \times \mathcal{T}^r(v)$.

Proof. Without loss of generality, suppose that v is either strictly to the right of L or on L. Note that each point $v_i \in V\backslash\{v\}$ corresponds to exactly one outer tangent line $T^r(v_i|v)$, thereby exactly one breakpoint $t_L^r(v_i|v)$. Such one-to-one correspondence can be easily done in $O(1)$ time. Therefore, equivalently we are computing sequences of points in $V\backslash\{v\}$, instead of breakpoints.

In the following, we consider two cases about the relative position between L and $C_\gamma(v)$, (1) L intersects with $C_\gamma(v)$ at zero or one point, (2) L intersects with $C_\gamma(v)$ at two points.

Case (1): Let θ_L be the angle of the upward direction along L. See Figure 4(a). We classify the points in $V\backslash\{v\}$ by their polar angles with respect to v. Let $P_1(v)$ denote the sequence of those points with polar angles in the interval $(\theta_L, \theta_L + \pi)$ and sorted in CCW order. Similarly, let $P_2(v)$ be the sequence of points with polar angles in $(\theta_L + \pi, \theta_L)$ and sorted in CCW order. Obviously, $P_1(v)$ and $P_2(v)$ together satisfy condition (c). (Note that points with polar angles θ_L and $\theta_L + \pi$ are ignored, since they correspond to outer tangent lines parallel to L.)

By general position assumption, we can observe that, for any two distinct points v_i, v_j in $P_1(v)$, $t_L^r(v_i|v)$ is strictly above $t_L^r(v_j|v)$ if and only if v_i precedes v_j in $P_1(v)$. Thus, the ordering of points in $P_1(v)$ implicitly describes an ordering of their corresponding breakpoints in decreasing y-coordinates. Similarly, the ordering in $P_2(v)$ implies an ordering of corresponding breakpoints in decreasing y-coordinates. It follows that both $P_1(v)$ and $P_2(v)$ satisfy condition (b).

As for condition (a), both $P_1(v)$ and $P_2(v)$ are of length $O(n)$ by definition. Also, since we have pre-computed the sequence $P(v)$ as all points in $V\backslash\{v\}$

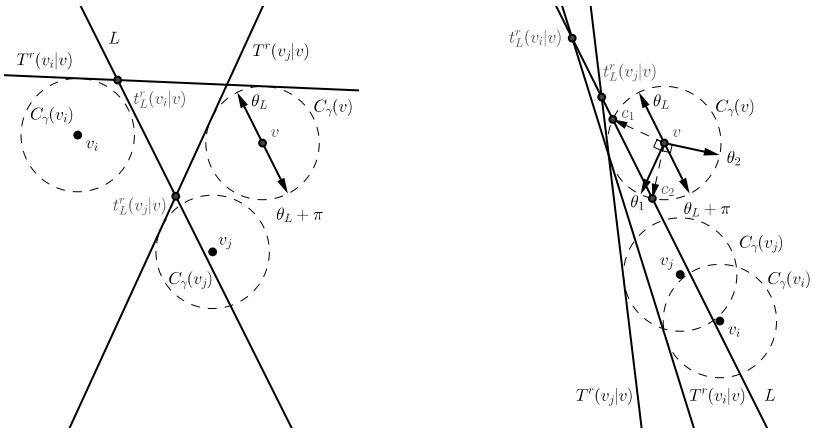

(a) no intersection (b) two intersection points

Fig. 4. Two subcases about how $C_\gamma(v)$ intersects C.

sorted in CCW order, $P_1(v)$ and $P_2(v)$ can be implicitly represented as concatenations of subsequences of $P(v)$. This can be done in $O(\log n)$ time by searching in $P(v)$ the foremost elements with polar angles larger than θ_L and $\theta_L + \pi$, respectively.

Case (2): Suppose that the two intersection points between L and $C_\gamma(v)$ are c_1 and c_2, where c_1 is above c_2. Let $\theta_1 = \theta_1' + \pi/2$ and $\theta_2 = \theta_2' + \pi/2$, in which θ_1' and θ_2' are respectively the polar angles of c_1 and c_2 with respect to v. See Figure 4(b). By assumption, we have that $\theta_L < \theta_1' < \theta_L + \pi/2 < \theta_2' < \theta_L + \pi$, which implies that $\theta_1 \in (\theta_L, \theta_L + \pi)$ and $\theta_2 \in (\theta_L + \pi, \theta_L)$.

We divide the points in $V \backslash \{v\}$ into four sequences $P_1(v)$, $P_2(v)$, $P_3(v)$, and $P_4(v)$ by their polar angles with respect to v. $P_1(v)$ consists of points with polar angles in (θ_L, θ_1), $P_2(v)$ in $[\theta_1, \theta_L + \pi)$, $P_3(v)$ in $(\theta_L + \pi, \theta_2]$, and $P_4(v)$ in (θ_2, θ_L), all sorted in CCW order. It follows that the four sequences satisfy conditions (c).

Condition (a) and (b) hold for $P_1(v)$ and $P_4(v)$ from similar discussion as above. However, for any two distinct points v_i, v_j in $P_2(v)$, we can observe that $t_L^r(v_i|v)$ is strictly below $t_L^r(v_j|v)$ if and only if v_i precedes v_j in $P_2(v)$. Similarly, the argument holds for $P_3(v)$. Thus, what satisfy condition (b) are actually the reverse sequences of $P_2(v)$ and $P_3(v)$, which can also be obtained in $O(\log n)$ time, satisfying condition (a). □

By Lemma 13.(c), searching in $L \times \mathcal{T}^r(v)$ is equivalent to searching in the $O(1)$ sequences of breakpoints, which can be computed more efficiently than the obvious way. Besides, we can also obtain a symmetrical lemma constructing sequences for $L \times \mathcal{T}^l(v)$. In the following, we show how to perform a binary search within these sequences.

Lemma 14 *With an $O(n^2 \log n)$-time preprocessing, given an arbitrary line L, a local optimal point x_L^* can be computed in $O(n \log^2 n)$ time.*

Proof. By Lemma 12, the searching of x_L^* can be done within $L \times \mathcal{T}$ and $L \times \mathcal{C}(V)$. $L \times \mathcal{T}$ can be further divided into $L \times \mathcal{T}^r(v)$ and $L \times \mathcal{T}^l(v)$ for each $v \in V$. By Lemma 13, these $2n$ sets can be replaced by $O(n)$ sorted sequences of breakpoints on L. Besides, $L \times \mathcal{C}(V)$ consists of no more than $2n$ breakpoints, which can be computed and arranged into a sorted sequence in decreasing y-coordinates. Therefore, we can construct $N_0 = O(n)$ sequences $P_1, P_2, \cdots, P_{N_0}$ of breakpoints, each of length $O(n)$ and sorted in decreasing y-coordinates.

The searching in the N_0 sorted sequences is done by performing parametric search for parallel binary searches, introduced in [1]. The technique we used here is similar to the algorithm in [1], but uses a different weighting scheme. For each sorted sequence P_j, $1 \leq j \leq N_0$, we first obtain its middle element x_j, and associate x_j with a weight m_j equal to the number of elements in P_j. Then, we compute the *weighted median* [18] of the N_0 middle elements, defined as the element x such that $\sum\{m_j | x_j \text{ is above } x\} \geq \sum m_j/2$ and $\sum\{m_j | x_j \text{ is below } x\} \geq \sum m_j/2$. Finally, we apply Lemma 10 on the point x. If x is a strong $(1|1)_R$-centroid, of course it is local optimal. If not, Assumption 1 holds and $Wedge(x)$ can be computed. If $Wedge(x)$ is sideward, a local optimal point $x_L^* = x$ is directly found. Otherwise, $Wedge(x)$ is either upward or downward, and thus all breakpoints on the opposite half can be pruned by Lemma 9. The pruning makes a portion of sequences, that possesses over half of total breakpoints by the definition of weighted median, lose at least a quarter of their elements. Hence, at least one-eighths of breakpoints are pruned. By repeating the above process, we can find x_L^* in at most $O(\log n)$ iterations.

The time complexity for finding x_L^* is analyzed as follows. By Lemma 13, constructing sorted sequences for $L \times \mathcal{T}^r(v)$ and $L \times \mathcal{T}^l(v)$ for all $v \in V$ takes $O(n \log n)$ time. Computing and sorting $L \times \mathcal{C}(V)$ also takes $O(n \log n)$ time. There are at most $O(\log n)$ iterations of the pruning process. At each iteration, the N_0 middle elements and their weighted median x can be obtained in $O(N_0) = O(n)$ time by the linear-time weighted selection algorithm [18]. Then, the computation of $Wedge(x)$ takes $O(n \log n)$ time by Lemma 10. Finally, the pruning of those sequences can be done in $O(n)$ time. In summary, the searching of x_L^* requires $O(n \log n) + O(\log n) \times O(n \log n) = O(n \log^2 n)$ time. □

We remark that, by Lemma 14, it is easy to obtain an intermediate result for the $(1|1)_R$-centroid problem on the plane. By Lemma 5, there exists a $(1|1)_R$-centroid in $\mathcal{T} \times \mathcal{T}$, $\mathcal{T} \times \mathcal{C}(V)$, and $\mathcal{C}(V) \times \mathcal{C}(V)$. By applying Lemma 14 to the $O(n^2)$ lines in $\mathcal{T}$, the local optimum among the intersection points in $\mathcal{T} \times \mathcal{T}$ and $\mathcal{T} \times \mathcal{C}(V)$ can be obtained in $O(n^3 \log^2 n)$ time. By applying Theorem 1 on the $O(n^2)$ intersection points in $\mathcal{C}(V) \times \mathcal{C}(V)$, the local optimum among them can be obtained in $O(n^3 \log n)$ time. Thus, we can find a $(1|1)_R$-centroid in $O(n^3 \log^2 n)$ time, a nearly $O(n^2)$ improvement over the $O(n^5 \log n)$ bound in [5].

4 $(1|1)_R$-Centroid on the Plane

In this section, we study the $(1|1)_R$-centroid problem and propose an improved algorithm of time complexity $O(n^2 \log n)$. This algorithm is as efficient as the

best-so-far algorithm for the (1|1)-centroid problem, but based on a completely different approach.

In Subsection 4.1, we extend the algorithm of Lemma 14 to develop a procedure allowing us to prune candidate points with respect to a given vertical line. Then, in Subsection 4.2, we show how to compute a $(1|1)_R$-centroid in $O(n^2 \log n)$ time based on this newly-developed pruning procedure.

4.1 Pruning with Respect to a Vertical Line

Let L be an arbitrary vertical line on the plane. We call the half-plane strictly to the left of L the *left plane* of L and the one strictly to its right the *right plane* of L. A sideward wedge of some point on L is said to be *rightward* (*leftward*) if it intersects the right (left) plane of L. We can observe that, if there is some point $x \in L$ such that $Wedge(x)$ is rightward, every point x' on the left plane of L can be pruned, since $W^*(x') \geq W^*(x)$ by Lemma 9. Similarly, if $Wedge(x)$ is leftward, points on the right plane of L can be pruned. Although the power of wedges is not fully exerted in this way, pruning via vertical lines and sideward wedges is superior than directly via wedges due to predictable pruning regions.

Therefore, in this subsection we describe how to design a procedure that enables us to prune either the left or the right plane of a given vertical line L. As mentioned above, the key point is the searching of sideward wedges on L. It is achieved by carrying out three conditional phases. In the first phase, we try to find some proper breakpoints with sideward wedges. If failed, we pick some representative point in the second phase and check its wedge to determine whether or not sideward wedges exist. Finally, in case of their nonexistence, we show that their functional alternative can be computed, called the *pseudo wedge*, that still allows us to prune the left or right plane of L. In the following, we develop a series of lemmas to demonstrate the details of the three phases.

Property 1 *Given a point $x \in L$, for each possible direction of $Wedge(x)$, the corresponding $CA(x)$ satisfies the following conditions:*

Upward – $CA(x) \subseteq (0, \pi)$,
Downward – $CA(x) \subseteq (\pi, 2\pi)$,
Rightward – $0 \in CA(x)$,
Leftward – $\pi \in CA(x)$.

Proof. When $Wedge(x)$ is upward, by definition the beginning angle θ_b and the ending angle θ_e of $CA(x)$ must satisfy that both half-planes $H(F(\theta_b|x), y(\theta_b|x))$ and $H(F(\theta_e|x), y(\theta_e|x))$ include the half-line of L above x. It follows that $0 \leq \theta_b, \theta_e \leq \pi$, and thus $CA(x) \subseteq (0, \pi)$. (Recall that $\theta_b, \theta_e \notin CA(x)$.) The case that $Wedge(x)$ is downward can be proved in a symmetric way.

When $Wedge(x)$ is rightward, we can see that $H(F(\theta_b|x), y(\theta_b|x))$ must not contain the half-line of L above x, and thus $\pi < \theta_b < 2\pi$. By similar arguments, $0 < \theta_e < \pi$. Therefore, counterclockwise covering angles from θ_b to θ_e, $CA(x)$ must include the angle 0. The case that $Wedge(x)$ is leftward can be symmetrically proved. □

Lemma 15 *Let x_1, x_2 be two points on L, where x_1 is strictly above x_2. For any angle $0 \leq \theta \leq \pi$, $W(\theta|x_1) \leq W(\theta|x_2)$. Symmetrically, for $\pi \leq \theta \leq 2\pi$, $W(\theta|x_2) \leq W(\theta|x_1)$.*

Proof. For any angle $0 \leq \theta \leq \pi$, we can observe that $H^-(B(\theta|x_1), y(\theta|x_1)) \subset H^-(B(\theta|x_2), y(\theta|x_2))$, since x_1 is strictly above x_2. It follows that $W(\theta|x_1) \leq W(\theta|x_2)$. The second claim also holds by symmetric arguments. □

Lemma 16 *Let x be an arbitrary point on L. If $Wedge(x)$ is either upward or downward, for any point $x' \in L \backslash Wedge(x)$, $Wedge(x')$ has the same direction as $Wedge(x)$.*

Proof. By symmetry, we prove that, if $Wedge(x)$ is upward, $Wedge(x')$ is also upward for every $x' \in L$ strictly below x. By Property 1, the fact that $Wedge(x)$ is upward means that $CA(x) \subset (0, \pi)$ and thus $MA(x) \subset (0, \pi)$. Let x' be a point on L strictly below x. By Lemma 15, we have that $W(\theta|x') \geq W(\theta|x)$ for $0 < \theta < \pi$ and $W(\theta|x') \leq W(\theta|x)$ for $\pi \leq \theta \leq 2\pi$. It follows that $MA(x') \subset (0, \pi)$ and $CA(x') \subset (0, \pi)$, so $Wedge(x')$ is upward as well. □

Following from this lemma, if there exist two arbitrary points x_1 and x_2 on L with their wedges downward and upward, respectively, we can derive that x_1 must be strictly above x_2, and that points with sideward wedges or even strong $(1|1)_R$-centroids can locate only between x_1 and x_2. Thus, we can find sideward wedges between some specified downward and upward wedges. Let x_D be the lowermost breakpoint on L with its wedge downward, x_U the uppermost breakpoint on L with its wedge upward, and G_{DU} the open segment $\overline{x_D x_U} \backslash \{x_D, x_U\}$. (For ease of discussion, we assume that both x_D and x_U exist on L, and show how to resolve this assumption later by constructing a bounded box.) Again, x_D is strictly above x_U. Also, we have the following corollary by their definitions.

Corollary 17 *If there exist breakpoints in the segment G_{DU}, for any such breakpoint x, either x is a strong $(1|1)_R$-centroid or $Wedge(x)$ is sideward.*

Given x_D and x_U, the first phase can thus be done by checking whether there exist breakpoints in G_{DU} and picking any of them if exist. Supposing that the picked one is not a strong $(1|1)_R$-centroid, a sideward wedge is found by Corollary 17 and can be used for pruning. Notice that, when there are two or more such breakpoints, one may question whether their wedges are of the same direction, as different directions result in inconsistent pruning results. The following lemma answers the question in the positive.

Lemma 18 *Let x_1, x_2 be two distinct points on L, where x_1 is strictly above x_2 and none of them is a strong $(1|1)_R$-centroid. If $Wedge(x_1)$ and $Wedge(x_2)$ are both sideward, they are either both rightward or both leftward.*

Proof. We prove this lemma by contradiction. By symmetry, suppose the case that $Wedge(x_1)$ is rightward and $Wedge(x_2)$ is leftward. This case can be further divided into two subcases by whether or not $CA(x_1)$ and $CA(x_2)$ intersect.

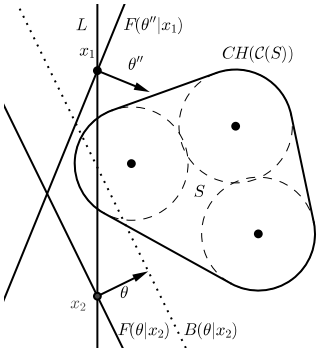

Fig. 5. $CH(\mathcal{C}(S))$ intersects L between x_1 and x_2.

Consider first that $CA(x_1)$ does not intersect $CA(x_2)$. Because $Wedge(x_1)$ is rightward, $0 \in CA(x_1)$ by Property 1. Thus, there exists an angle θ, $0 < \theta < \pi$, such that $\theta \in MA(x_1)$. Since x_1 is strictly above x_2, by Lemma 15 we have that $W^*(x_1) = W(\theta|x_1) \leq W(\theta|x_2) \leq W^*(x_2)$. Furthermore, since $Wedge(x_2)$ is leftward, we can see that $x_1 \notin Wedge(x_2)$ and therefore $W^*(x_1) \geq W^*(x_2)$ by Lemma 9. It follows that $W(\theta|x_2) = W^*(x_2)$ and thus $\theta \in MA(x_2)$. By definition, $MA(x_1) \subseteq CA(x_1)$ and $MA(x_2) \subseteq CA(x_2)$, which implies that $CA(x_1)$ and $CA(x_2)$ intersect at θ, contradicting the subcase assumption.

When $CA(x_1)$ intersects $CA(x_2)$, their intersection must be completely included in either $(0, \pi)$ or $(\pi, 2\pi)$ due to Assumption 1. By symmetry, we assume the latter subcase. Using similar arguments as above, we can find an angle θ', where $0 < \theta' < \pi$, such that $\theta' \in MA(x_1)$ and $\theta' \in MA(x_2)$. This is a contradiction, since $\theta' \notin (\pi, 2\pi)$.

Since both subcases do not hold, the lemma is proved. □

The second phase deals with the case that no breakpoint exists between x_D and x_U by determining the wedge direction of an arbitrary inner point in G_{DU}. We begin with several auxiliary lemmas.

Lemma 19 *Let x_1, x_2 be two distinct points on L such that $W^*(x_1) = W^*(x_2)$ and x_1 is strictly above x_2. There exists at least one breakpoint in the segment*

(a) $\overline{x_1x_2}\backslash\{x_2\}$, if $MA(x_2)$ intersects $(0, \pi)$ but $MA(x_1)$ does not,
(b) $\overline{x_1x_2}\backslash\{x_1\}$, if $MA(x_1)$ intersects $(\pi, 2\pi)$ but $MA(x_2)$ does not.

Proof. By symmetry, we only show the correctness of condition (a). From its assumption, there exists an angle θ, where $0 < \theta < \pi$, such that $\theta \in MA(x_2)$. Let $S = V \bigcap H^-(B(\theta|x_2), y(\theta|x_2))$. By definition, we have that $W(S) = W(\theta|x_2) = W^*(x_2) = W^*(x_1)$ and $CH(\mathcal{C}(S)) \subset H^-(F(\theta|x_2), y(\theta|x_2))$, which implies that $CH(\mathcal{C}(S))$ is strictly above $F(\theta|x_2)$. (See Figure 5.)

We first claim that $CH(\mathcal{C}(S))$ intersects L. If not, there must exist an angle θ', where $0 < \theta' < \pi$, such that $CH(\mathcal{C}(S)) \subset H^-(F(\theta'|x_1), y(\theta'|x_1))$, that is,

$S \subset H^-(B(\theta'|x_1), y(\theta'|x_1))$. By definition, $W^*(x_1) \geq W(\theta'|x_1) \geq W(S)$. Since $W^*(x_1) = W(S)$, $W(\theta'|x_1) = W^*(x_1)$ and thus $\theta' \in MA(x_1)$, which contradicts the condition that $MA(x_1)$ does not intersect $(0, \pi)$. Thus, the claim holds.

When $CH(\mathcal{C}(S))$ intersects L, x_1 locates either inside or outside $CH(\mathcal{C}(S))$. Since x_2 locates outside $CH(\mathcal{C}(S))$, in the former case the boundary of $CH(\mathcal{C}(S))$ intersects $\overline{x_1x_2}\backslash\{x_2\}$ and forms a breakpoint, thereby proves condition (a). On the other hand, if x_1 is outside $CH(\mathcal{C}(S))$, again there exists an angle θ'' such that $CH(\mathcal{C}(S)) \subset H^-(F(\theta''|x_1), y(\theta''|x_1))$. By similar arguments, we can show that $\theta'' \in MA(x_1)$. By assumption, θ'' must belong to $(\pi, 2\pi)$, which implies that $CH(\mathcal{C}(S))$ is strictly below $F(\theta''|x_1)$. Since $CH(\mathcal{C}(S))$ is strictly above $F(\theta|x_2)$ as mentioned, any intersection point between $CH(\mathcal{C}(S))$ and L should be inner to $\overline{x_1x_2}$. Therefore, the lemma holds. □

Lemma 20 *Let G be a line segment connecting two consecutive breakpoints on L. For any two distinct points x_1, x_2 inner to G, $W^*(x_1) = W^*(x_2)$.*

Proof. Suppose to the contrary that $W^*(x_1) \neq W^*(x_2)$. By Lemma 11, there exists at least one breakpoint in $\overline{x_1x_2}$, which contradicts the definition of G. Thus, the lemma holds. □

Lemma 21 *When there is no breakpoint between x_D and x_U, any two distinct points x_1, x_2 in G_{DU} have the same wedge direction, if they are not strong $(1|1)_R$-centroids.*

Proof. Suppose by contradiction that the directions of their wedges are different. By Lemmas 16 and 18, there are only two possible cases.

(1) $Wedge(x_1)$ is downward, and $Wedge(x_2)$ is either sideward or upward.
(2) $Wedge(x_1)$ is sideward, and $Wedge(x_2)$ is upward.

In the following, we show that both cases do not hold.

Case (1): Because $Wedge(x_1)$ is downward, we have that $CA(x_1) \subseteq (\pi, 2\pi)$ by Property 1 and thus $MA(x_1)$ does not intersect $(0, \pi)$. On the other hand, whether $Wedge(x_2)$ is sideward or upward, we can see that $CA(x_2)$ and $MA(x_2)$ intersect $(0, \pi)$ by again Property 1. Since $W^*(x_1) = W^*(x_2)$ by Lemma 20, the status of the two points satisfies the condition (a) of Lemma 19, so at least one breakpoint exists between x_1 and x_2. By definitions of x_1 and x_2, this breakpoint is inner to G_{DU}, thereby contradicts the assumption. Therefore, Case (1) does not hold.

Case (2): The proof of Case (2) is symmetric to that of Case (1). The condition (b) of Lemma 19 can be applied similarly to show the existence of at least one breakpoint between x_1 and x_2, again a contradiction.

Combining the above discussions, we prove that the wedges of x_1 and x_2 are of the same direction, thereby completes the proof of this lemma. □

This lemma enables us to pick an arbitrary point in G_{DU}, e.g., the bisector point x_B of x_D and x_U, as the representative of all inner points in G_{DU}. If x_B is not a strong $(1|1)_R$-centroid and $Wedge(x_B)$ is sideward, the second phase finishes with a sideward wedge found. Otherwise, if $Wedge(x_B)$ is downward or upward, we can derive the following and have to invoke the third phase.

Lemma 22 *If there is no breakpoint between x_D and x_U and $Wedge(x_B)$ is not sideward, there exist neither strong $(1|1)_R$-centroids nor points with sideward wedges on L.*

Proof. By Lemma 16, this lemma holds for points not in G_{DU}. Without loss of generality, suppose that $Wedge(x_B)$ is downward. For all points in G_{DU} above x_B, the lemma holds by again Lemma 16.

Consider an arbitrary point $x \in \overline{x_B x_U} \backslash \{x_B, x_U\}$. We first show that x is not a strong $(1|1)_R$-centroid. Suppose to the contrary that x really is. By definition, we have that $\delta(CA(x)) > \pi$, and thus $CA(x)$ and $MA(x)$ intersect $(0, \pi)$. On the other hand, $CA(x_B)$ and $MA(x_B)$ do not intersect $(0, \pi)$ due to downward $Wedge(x_B)$ and Property 1. Since $W^*(x_B) = W^*(x)$ by Lemma 20, applying the condition (a) of Lemma 19 to x_B and x shows that at least one breakpoint exists between them, which contradicts the no-breakpoint assumption. Now that x is not a strong $(1|1)_R$-centroid, it must have a downward wedge, as x_B does by Lemma 21. Therefore, the lemma holds for all points on L. □

When L satisfies Lemma 22, it consists of only points with downward or upward wedges, and is said to be *non-leaning*. Obviously, our pruning strategy via sideward wedges could not apply to such non-leaning lines. The third phase overcomes this obstacle by constructing a functional alternative of sideward wedges, called the pseudo wedge, on either x_D or x_U, so that pruning with respect to L is still achievable. Again, we start with auxiliary lemmas.

Lemma 23 *If L is non-leaning, the following statements hold:*

(a) $W^*(x_D) \neq W^*(x_U)$,
(b) $W^*(x) = \max\{W^*(x_D), W^*(x_U)\}$ *for all points* $x \in G_{DU}$.

Proof. We prove the correctness of statement (a) by contradiction, and suppose that $W^*(x_D) = W^*(x_U)$. Besides, the fact that L is non-leaning implies that no breakpoint exists in G_{DU}. By Lemmas 22 and 21, the wedges of all points in G_{DU} are of the same direction, either downward or upward. Suppose the downward case by symmetry, and pick an arbitrary point in G_{DU}, say, x_B. Since $Wedge(x_B)$ is downward, we have that $MA(x_B)$ does not intersect $(0, \pi)$. Oppositely, by definition $Wedge(x_U)$ is upward, so $CA(x_U)$ and $MA(x_U)$ are included in $(0, \pi)$. Because x_B is strictly above x_U and $W^*(x_D) = W^*(x_U)$, according to the condition (a) of Lemma 19, there exists at least one breakpoint in $\overline{x_B x_U} \backslash \{x_U\}$, which is a contradiction. Therefore, statement (a) holds.

The proof of statement (b) is also done by contradiction. By symmetry, assume that $W^*(x_D) > W^*(x_U)$ in statement (a). Consider an arbitrary point

$x \in G_{DU}$. By Lemma 7, we have that $W^*(x) \leq \max\{W^*(x_D), W^*(x_U)\} = W^*(x_D)$. Suppose that the equality does not hold. Then, by Lemma 11, at least one breakpoint exists in the segment $\overline{x_D x}\backslash\{x_D\}$, contradicting the no-breakpoint fact. Thus, $W^*(x) = W^*(x_D)$ and statement (b) holds. □

Let $W_1 = \max\{W^*(x_D), W^*(x_U)\}$. We are going to define the pseudo wedge on either x_U or x_D, depending on which one has the smaller weight loss. We consider first the case that $W^*(x_D) > W^*(x_U)$, and obtain the following.

Lemma 24 *If L is non-leaning and $W^*(x_D) > W^*(x_U)$, there exists one angle θ for x_U, where $\pi \leq \theta \leq 2\pi$, such that $W(H(B(\theta|x_U), y(\theta|x_U))) \geq W_1$.*

Proof. We first show that there exists at least a subset $S \subseteq V$ with $W(S) = W_1$, such that x_U locates on the upper boundary of $CH(\mathcal{C}(S))$. Let x be the point strictly above but arbitrarily close to x_U on L. By Lemma 23, $W^*(x) = W_1$, hence $W^*(x) > W^*(x_U)$ by case assumption. It follows that $x_U \in Wedge(x)$ by Lemma 9 and $Wedge(x)$ must be downward. By Property 1, we have that $CA(x) \subseteq (\pi, 2\pi)$. Thus, there exists an angle $\theta' \in MA(x)$, where $\pi < \theta' < 2\pi$, such that $W(H^-(B(\theta'|x), y(\theta'|x))) = W(\theta'|x) = W_1$.

Let $S = V \bigcap H^-(B(\theta'|x), y(\theta'|x))$. Since $W^*(x_U) < W_1 = W(S)$, x_U is inside $CH(\mathcal{C}(S))$ by Lemma 2. Oppositely, by the definition of S, x is outside the convex hull $CH(\mathcal{C}(S))$. It implies that x_U is the topmost intersection point between $CH(\mathcal{C}(S))$ and L, hence on the upper boundary of $CH(\mathcal{C}(S))$. (It is possible that x_U locates at the leftmost or the rightmost point of $CH(\mathcal{C}(S))$.)

The claimed angle θ is obtained as follows. Since x_U is a boundary point of $CH(\mathcal{C}(S))$, there exists a line F passing through x_U and tangent to $CH(\mathcal{C}(S))$. Let θ be the angle satisfying that $F(\theta|x_U) = F$, $\pi \leq \theta \leq 2\pi$, and $CH(\mathcal{C}(S)) \subset H(F(\theta|x_U), y(\theta|x_U))$. Obviously, we have that $S \subset H(B(\theta|x_U), y(\theta|x_U))$ and thus $W(H(B(\theta|x_U), y(\theta|x_U))) \geq W^*(S) = W_1$. □

Let θ_U be an arbitrary angle satisfying the conditions of Lemma 24. We apply the line $F(\theta_U|x_U)$ for trimming the region of $Wedge(x_U)$, so that a sideward wedge can be obtained. Let $PW(x_U)$, called the *pseudo wedge* of x_U, denote the intersection of $Wedge(x_U)$ and $H(F(\theta_U|x_U), y(\theta_U|x_U))$. Deriving from the three facts that $Wedge(x_U)$ is upward, $\delta(Wedge(x_U)) < \pi$, and $\pi \leq \theta_U \leq 2\pi$, we can observe that either $PW(x_U)$ is x_U itself, or it intersects only one of the right and left plane of L. In the two circumstances, $PW(x_U)$ is said to be *null* or *sideward*, respectively. The pseudo wedge has similar functionality as wedges, as shown in the following corollary.

Corollary 25 *For any point $x' \notin PW(x_U)$, $W^*(x') \geq W^*(x_U)$.*

Proof. If $x' \notin Wedge(x_U)$, the lemma directly holds by Lemma 9. Otherwise, we have that $x' \notin H(F(\theta|x_U), y(\theta|x_U))$ and thus $H(B(\theta|x'), y(\theta|x'))$ contains $H(B(\theta|x_U), y(\theta|x_U))$. Then, by Lemma 24, $W^*(x') \geq W(H(B(\theta|x'), y(\theta|x'))) \geq W(H(B(\theta|x_U), y(\theta|x_U))) \geq W_1$, thereby completes the proof. □

By this lemma, if $PW(x_U)$ is found to be sideward, points on the opposite half-plane with respect to L can be pruned. If $PW(x_U)$ is null, x_U becomes another kind of strong $(1|1)_R$-centroids, in the meaning that it is also an immediate solution to the $(1|1)_R$-centroid problem. Without confusion, we call x_U a *conditional* $(1|1)_R$*-centroid* in the latter case.

On the other hand, considering the reverse case that $W^*(x_D) < W^*(x_U)$, we can also obtain an angle x_D and a pseudo wedge $PW(x_D)$ for x_D by symmetric arguments. Then, either $PW(x_D)$ is sideward and the opposite side of L can be pruned, or $PW(x_D)$ itself is a conditional $(1|1)_R$-centroid. Thus, the third phase solves the problem of the nonexistence of sideward wedges.

Recall that the three phases of searching sideward wedges is based on the existence of x_D and x_U on L, which was not guaranteed before. Here we show that, by constructing appropriate border lines, we can guarantee the existence of x_D and x_U while searching between these border lines. The *bounding box* is defined as the smallest axis-aligned rectangle that encloses all circles in $\mathcal{C}(V)$. Obviously, any point x outside the box satisfies that $W^*(x) = W(V)$ and must not be a $(1|1)_R$-centroid. Thus, given a vertical line not intersecting the box, the half-plane to be pruned is trivially decided. Moreover, let T_{top} and T_{btm} be two arbitrary horizontal lines strictly above and below the bounding box, respectively. We can obtain the following.

Lemma 26 *Let L be an arbitrary vertical line intersecting the bounding box, and x'_D and x'_U denote its intersection points with T_{top} and T_{btm}, respectively. $Wedge(x'_D)$ is downward and $Wedge(x'_U)$ is upward.*

Proof. Consider the case about $Wedge(x'_D)$. As described above, we know the fact that $W^*(x_D) = W(V)$. Let θ be an arbitrary angle with $0 \leq \theta \leq \pi$. We can observe that $H^-(F(\theta|x'_D), y(\theta|x'_D))$ cannot contain all circles in $\mathcal{C}(V)$, that is, $V \not\subset H^-(B(\theta|x'_D), y(\theta|x'_D))$. This implies that $W(\theta|x'_D) < W^*(x'_D)$ and $\theta \notin MA(x'_D)$. Therefore, we have that $MA(x'_D) \subset (\pi, 2\pi)$ and $Wedge(x'_D)$ is downward by Property 1. By similar arguments, we can show that $Wedge(x'_U)$ is upward. Thus, the lemma holds. □

According to this lemma, by inserting T_{top} and T_{btm} into $\mathcal{T}$, the existence of x_D and x_U is enforced for any vertical line intersecting the bounding box. Besides, it is obvious to see that the insertion does not affect the correctness of all lemmas developed so far.

Summarizing the above discussion, the whole picture of our desired pruning procedure can be described as follows. In the beginning, we perform a preprocessing to obtain the bounding box and then add T_{top} and T_{btm} into $\mathcal{T}$. Now, given a vertical line L, whether to prune its left or right plane can be determined by the following steps.

1. If L does not intersect the bounding box, prune the half-plane not containing the box.
2. Compute x_D and x_U on L.

3. Find a sideward wedge or pseudo wedge via three forementioned phases. (Terminate whenever a strong or conditional $(1|1)_R$-centroid is found.)
   (a) If breakpoints exist between x_D and x_U, pick any of them and check it.
   (b) If no such breakpoint, decide whether L is non-leaning by checking x_B.
   (c) If L is non-leaning, compute $PW(x_U)$ or $PW(x_D)$ depending on which of x_U and x_D has smaller weight loss.
4. Prune the right or left plane of L according to the direction of the sideward wedge or pseudo wedge.

The correctness of this procedure follows from the developed lemmas. Any vertical line not intersecting the bounding box is trivially dealt with in Step 1, due to the property of the box. When L intersects the box, by Lemma 26, x_D and x_U can certainly be found in Step 2. The three sub-steps of Step 3 correspond to the three searching phases. When L is not non-leaning, a sideward wedge is found, either at some breakpoint between x_D and x_U in Step 3(a) by Corollary 17, or at x_B in Step 3(b) by Lemma 21. Otherwise, according to Lemma 24 or its symmetric version, a pseudo wedge can be built in Step 3(c) for x_U or x_D, respectively. Finally in Step 4, whether to prune the left or right plane of L can be determined via the just-found sideward wedge or pseudo wedge, by respectively Lemma 9 or Corollary 25.

The time complexity of this procedure is analyzed as follows. The preprocessing for computing the bounding box trivially takes $O(n)$ time. In Step 1, any vertical line not intersecting the box can be identified and dealt with in $O(1)$ time. Finding x_D and x_U in Step 2 requires the help of the binary-search algorithm developed in 3.2. Although the algorithm is designed to find a local optimal point, we can easily observe that slightly modifying its objective makes it applicable to this purpose without changing its time complexity. Thus, Step 2 can be done in $O(n \log^2 n)$ time by Lemma 14.

In Step 3(a), all breakpoints between x_D and x_U can be found in $O(n \log n)$ time as follows. As done in Lemma 14, we first list all breakpoints on L by $O(n)$ sorted sequences of length $O(n)$, which takes $O(n \log n)$ time. Then, by performing binary search with the y-coordinates of x_D and x_U, we can find within each sequence the breakpoints between them in $O(\log n)$ time. In Step 3(a) or 3(b), checking a picked point x is done by computing $CA(x)$, that requires $O(n \log n)$ time by Lemma 10. To compute the pseudo wedge in Step 3(c), the angle θ_U satisfying Lemma 24, or symmetrically θ_D, can be computed in $O(n \log n)$ time by sweeping technique as in Lemma 10. Thus, $PW(x_U)$ or $PW(x_D)$ can be computed in $O(n \log n)$ time. Finally, the pruning decision in Step 4 takes $O(1)$ time. Summarizing the above, these steps require $O(n \log^2 n)$ time in total. Since the invocation of Lemma 14 needs an additional $O(n^2 \log n)$-time preprocessing, we have the following result.

Lemma 27 *With an $O(n^2 \log n)$-time preprocessing, whether to prune the right or left plane of a given vertical line L can be determined in $O(n \log^2 n)$ time.*

4.2 Searching on the Euclidean Plane

In this subsection, we come back to the $(1|1)_R$-centroid problem. Recall that, by Lemma 5, at least one $(1|1)_R$-centroid can be found in the three sets of intersection points $\mathcal{T} \times \mathcal{T}$, $\mathcal{C}(V) \times \mathcal{T}$, and $\mathcal{C}(V) \times \mathcal{C}(V)$, which consist of total $O(n^4)$ points. Let $\mathcal{L}$ denote the set of all vertical lines passing through these $O(n^4)$ intersection points. By definition, there exists a vertical line $L^* \in \mathcal{L}$ such that its local optimal point is a $(1|1)_R$-centroid. Conceptually, with the help of Lemma 27, L^* can be derived by applying prune-and-search approach to $\mathcal{L}$: pick the vertical line L from $\mathcal{L}$ with median x-coordinates, determine by Lemma 27 whether the right or left plane of L should be pruned, discard lines of $\mathcal{L}$ in the pruned half-plane, and repeat above until two vertical lines left. Obviously, it costs too much if this approach is carried out by explicitly generating and sorting the $O(n^4)$ lines. However, by separately dealing with each of the three sets, we can implicitly maintain sorted sequences of these lines and apply the prune-and-search approach.

Let $\mathcal{L}_\mathcal{T}$, $\mathcal{L}_\mathcal{M}$, and $\mathcal{L}_\mathcal{C}$ be the sets of all vertical lines passing through the intersection points in $\mathcal{T} \times \mathcal{T}$, $\mathcal{C}(V) \times \mathcal{T}$, and $\mathcal{C}(V) \times \mathcal{C}(V)$, respectively. A *local optimal line* of $\mathcal{L}_\mathcal{T}$ is a vertical line L_t^* such that its local optimal point has weight loss no larger than those of points in $\mathcal{T} \times \mathcal{T}$. The local optimal lines L_m^* and L_c^* can be similarly defined for $\mathcal{L}_\mathcal{M}$ and $\mathcal{L}_\mathcal{C}$, respectively. We will adopt different prune-and-search techniques to find the local optimal lines in the three sets, as shown in the following lemmas.

Lemma 28 *A local optimal line L_t^* of $\mathcal{L}_\mathcal{T}$ can be found in $O(n^2 \log n)$ time.*

Proof. Let $N_1 = |\mathcal{T}|$. By definition, there are $(N_1)^2$ intersection points in $\mathcal{T} \times \mathcal{T}$ and $(N_1)^2$ vertical lines in $\mathcal{L}_\mathcal{T}$. For efficiently searching within these vertical lines, we apply the ingenious idea of parametric search via parallel sorting algorithms, proposed by Megiddo [15].

Consider two arbitrary lines $T_g, T_h \in \mathcal{T}$. If they are not parallel, let t_{gh} be their intersection point and L_{gh} be the vertical line passing through t_{gh}. Suppose that T_g is above T_h in the left plane of L_{gh}. If applying Lemma 27 to L_{gh} prunes its right plane, T_g is above T_h in the remained left plane. On the other hand, if the left plane of L_{gh} is pruned, T_h is above T_g in the remained right plane. Therefore, L_{gh} can be treated as a "*comparison*" between T_g and T_h, in the sense that applying Lemma 27 to L_{gh} determines their ordering in the remained half-plane. It also decides the ordering of their intersection points with the undetermined local optimal line L_t^*, since the pruning ensures that a local optimal line stays in the remained half-plane.

It follows that, by resolving comparisons, the process of pruning vertical lines in $\mathcal{L}_\mathcal{T}$ to find L_t^* can be reduced to the problem of determining the ordering of the intersection points of the N_1 lines with L_t^*, or say, the sorting of these intersection points on L_t^*. While resolving comparisons during the sorting process, we can simultaneously maintain the remained half-plane by two vertical lines as its boundaries. Thus, after resolving all comparisons in $\mathcal{L}_\mathcal{T}$, one of the two boundaries must be a local optimal line. As we know, the most efficient way to

obtain the ordering is to apply some optimal sorting algorithm A_S, which needs to resolve only $O(N_1 \log N_1)$ comparisons, instead of $(N_1)^2$ comparisons. Since resolving each comparison takes $O(n \log^2 n)$ time by Lemma 27, the sorting is done in $O(n \log^2 n) \times O(N_1 \log^2 N_1) = O(n^3 \log^3 n)$ time, so is the finding of L_t^*.

However, Megiddo [15] observed that, when multiple comparisons can be indirectly resolved in a batch, simulating parallel sorting algorithms in a sequential way naturally provides the scheme for batching comparisons, thereby outperform the case of applying A_S. Let A_P be an arbitrary cost-optimal parallel sorting algorithm that runs in $O(\log n)$ steps on $O(n)$ processors, e.g., the parallel merge sort in [2]. Using A_P to sort the N_1 lines in $\mathcal{L}_\mathcal{T}$ on L_t^* takes $O(\log N_1)$ parallel steps. At each parallel step, there are $k = O(N_1)$ comparisons $L_1, L_2, \cdots, L_k$ to be resolved. We select the one with median x-coordinate among them, which is supposed to be some L_i. If applying Lemma 27 to L_i prunes its left plane, for each comparison L_j to the left of L_i, the ordering of the corresponding lines of L_j in the remained right plane of L_i is directly known. Thus, the $O(k/2)$ comparisons to the left of L_i are indirectly resolved in $O(k/2)$ time. If otherwise the right plane of L_i is pruned, the $O(k/2)$ comparisons to its right are resolved in $O(k/2)$ time. By repeating this process of selecting medians and pruning on the remaining elements $O(\log k)$ times, all k comparisons can be resolved, which takes $O(n \log^2 n) \times O(\log k) + O(k + k/2 + k/4 + \cdots) = O(n \log^3 n) + O(N_1) = O(n^2)$ time. Therefore, going through $O(\log N_1)$ parallel steps of A_P requires $O(n^2 \log N_1) = O(n^2 \log n)$ time, which determines the ordering of lines in $\mathcal{L}_\mathcal{T}$ on L_t^* and also computes a local optimal line L_t^*. □

Lemma 29 *A local optimal line L_m^* of $\mathcal{L}_\mathcal{M}$ can be found in $O(n^2 \log n)$ time.*

Proof. To deal with the set $\mathcal{L}_\mathcal{M}$, we use the ideas similar to the proofs of Lemmas 13 and 14 in order to divide $\mathcal{C}(V) \times \mathcal{T}$ into sorted sequences of points. Given a fixed circle $C = C_\gamma(u_0)$ for some point $u_0 \in V$, we show that the intersection points in $C \times \mathcal{T}^r(v)$ and $C \times \mathcal{T}^l(v)$ for each $v \in V$ can be grouped into $O(1)$ sequences of length $O(n)$, which are sorted in increasing x-coordinates. Summarizing over all circles in $\mathcal{C}(V)$, there will be total $O(n^2)$ sequences of length $O(n)$, each of which maps to a sequence of $O(n)$ vertical lines sorted in increasing x-coordinates. Then, finding a local optimal line L_m^* can be done by performing prune-and-search to the $O(n^2)$ sequences of vertical lines via parallel binary searches. The details of these steps are described as follows.

First we discuss about the way for grouping intersection points in $C \times \mathcal{T}^r(v)$ and $C \times \mathcal{T}^l(v)$ for a fixed point $v \in V$, so that each of them can be represented by $O(1)$ subsequences of $P(v)$. By symmetry, only $C \times \mathcal{T}^r(v)$ is considered. Similar to Lemma 13, we are actually computing sequences of points in $V \backslash \{v\}$ corresponding to these intersection points. For each $v_i \in V \backslash \{v\}$, the outer tangent line $T^r(v_i|v)$ may intersect C at two, one, or zero point. Let $t^r_{C,1}(v_i|v)$ and $t^r_{C,2}(v_i|v)$ denote the first and second points, respectively, at which $T^r(v_i|v)$ intersects C along the direction from v to v_i. Note that, when $T^r(v_i|v)$ intersects C at less than two points, $t^r_{C,2}(v_i|v)$ or both of them will be *null*.

In the following, we consider the sequence computation under two cases about the relationship between u_0 and v, (1) $u_0 = v$, and (2) $u_0 \neq v$.

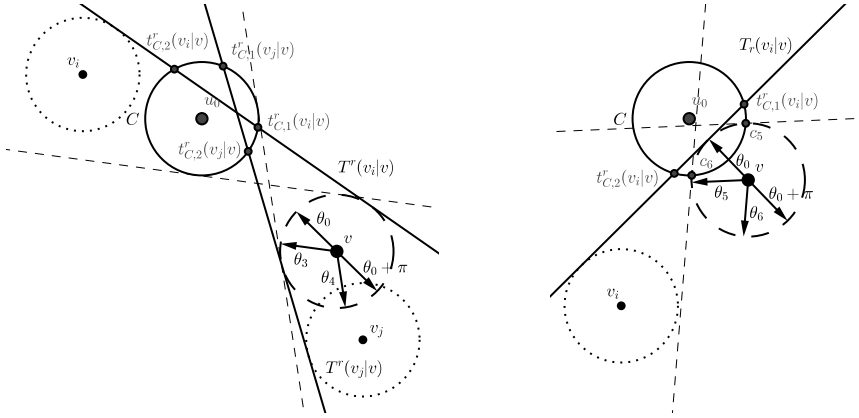

Fig. 6. Two subcases about how $C_\gamma(v)$ intersects C.

Case (1): Since v coincides with u_0, $C \times \mathcal{T}^r(v)$ is just the set of tangent points $t^r_{C,1}(v_i|v)$ for all $v_i \in V\backslash\{v\}$. It is easy to see the the angular sorted sequence $P(v)$ directly corresponds to a sorted sequence of these $n-1$ tangent points in CCW order. $P(v)$ can be further partitioned into two sub-sequences $P_1(v)$ and $P_2(v)$, which consist of points in $V\backslash\{v\}$ with polar angles (with respect to v) in the intervals $[0, \pi)$ and $[\pi, 2\pi)$, respectively. Since they are sorted in CCW order, we have that intersection points corresponding to $P_2(v)$ and to the reverse of $P_1(v)$ are sorted in increasing x-coordinates, as we required. Obviously, $P_1(v)$ and $P_2(v)$ are of length $O(n)$ and can be obtained in $O(\log n)$ time.

Case (2): Suppose without loss of generality that v locates on the lower left quadrant with respect to u_0, and let θ_0 be the polar angle of u_0 with respect to v. This case can be further divided into two subcases by whether or not $C_\gamma(v)$ intersects C at less than two points.

Consider first the subcase that they intersect at none or one point (see Figure 6(a).) Let θ_3 and θ_4 be the angles such that $T^r(y(\theta_3|v)|v)$ and $T^r(y(\theta_4|v)|v)$ are inner tangent to $C_\gamma(v)$ and C, where $\theta_3 \leq \theta_4$ (Note that $\theta_3 = \theta_4$ only when the two circles intersect at one point.) For each $v_i \in V\backslash\{v\}$, $T^r(v_i|v)$ does not intersect C, if the polar angle of v_i with respect to v is neither in $[\theta_0, \theta_3]$ nor in $[\theta_4, \theta_0 + \pi]$. We can implicitly obtain from $P(v)$ two subsequences $P_3(v)$ and $P_4(v)$, consisting of points with polar angles in $[\theta_0, \theta_3]$ and in $[\theta_4, \theta_0 + \pi]$, respectively. It can be observed that the sequence of points v_i listed in $P_3(v)$ corresponds to a sequence of intersection points $t^r_{C,1}(v_i|v)$ listed in clockwise (CW) order on C and, moreover, a sequence of $t^r_{C,2}(v_i|v)$ listed in CCW order on C. Symmetrically, the sequence of points v_j in $P_4(v)$ corresponds to a sequence of $t^r_{C,1}(v_j|v)$ in CCW order and a sequence of $t^r_{C,2}(v_j|v)$ in CW order. The four implicit sequences of intersection points on C can be further partitioned by a horizontal line L_h passing through

its center u_0, so that the resulted sequences are naturally sorted in either increasing or decreasing x-coordinates. Therefore, we can implicitly obtain at most eight sorted sequences of length $O(n)$ in replace of $C \times \mathcal{T}^r(v)$, by appropriately partitioning $P(v)$ in $O(\log n)$ time.
Consider that $C_\gamma(v)$ intersects $C_\gamma(u_0)$ at two points c_5 and c_6, where c_5 is to the upper right of c_6 (see Figure 6(b).) Let θ_5 and θ_6 be the angles such that $T^r(y(\theta_5|v)|v)$ and $T^r(y(\theta_6|v)|v)$ are tangent to $C_\gamma(v)$ at c_5 and c_6, respectively. Again, $P(v)$ can be implicitly partitioned into three subsequences $P_5(v)$, $P_6(v)$, and $P_7(v)$, which consists of points with polar angles in $[\theta_0, \theta_5)$, $[\theta_5, \theta_6)$, and $[\theta_6, \theta_0 + \pi]$, respectively. By similar observations, $P_5(v)$ corresponds to two sequences of intersection points listed in CW and CCW order, respectively, and $P_7(v)$ corresponds to two sequences listed in CCW and CW order, respectively. However, the sequence of points v_i in $P_6(v)$ corresponds to the sequences of $t^r_{C,1}(v_i|v)$ and $t^r_{C,2}(v_i|v)$ listed in both CCW order. These sequences can also be partitioned by L_h into sequences sorted in x-coordinates. It follows that we can implicitly obtain at most twelve sorted sequences of length $O(n)$ in replace of $C \times \mathcal{T}^r(v)$ in $O(\log n)$ time.

According to the above discussion, for any two points $u, v \in V$, $C_\gamma(u) \times \mathcal{T}^r(v)$ and $C_\gamma(u) \times \mathcal{T}^l(v)$ can be divided into $O(1)$ sequences in $O(\log n)$ time, each of which consists of $O(n)$ intersection points on $C_\gamma(u)$ sorted in increasing x-coordinates. Thus, $\mathcal{C}(V) \times \mathcal{T}$ can be re-organized as $O(n^2)$ sorted sequences of length $O(n)$ in $O(n^2 \log n)$ time, which correspond to $O(n^2)$ sorted sequences of $O(n)$ vertical lines. Now, we can perform parametric search for parallel binary search to these sequences of vertical lines, by similar techniques used in Lemma 14. For each of the $O(n^2)$ sequences, its middle element is first obtained and assigned with a weight equal to the sequence length in $O(1)$ time. Then, the weighted median L of these $O(n^2)$ elements are computed in $O(n^2)$ time [18]. By applying Lemma 27 to L in $O(n \log^2 n)$ time, at least one-eighths of total elements can be pruned from these sequences, taking another $O(n^2)$ time. Therefore, a single iteration of pruning requires $O(n^2)$ time. After $O(\log n)$ such iterations, a local optimal line L^*_m can be found in total $O(n^2 \log n)$ time, thereby proves the lemma. □

Lemma 30 *A local optimal line L^*_c of $\mathcal{L}_\mathcal{C}$ can be found in $O(n^2 \log n)$ time.*

Proof. There are at most $O(n^2)$ points in $\mathcal{C}(V) \times \mathcal{C}(V)$. Thus, $\mathcal{L}_\mathcal{C}$ can be obtained and sorted according to x-coordinates in $O(n^2 \log n)$ time. Then, by simply performing binary search with Lemma 27, a local optimal line L^*_c can be easily found in $O(\log n)$ iterations of pruning, which require total $O(n \log^3 n)$ time. In summary, the computation takes $O(n^2 \log n)$ time, and the lemma holds. □

By definition, L^* can be found among L^*_t, L^*_m, and L^*_c, which can be computed in $O(n^2 \log n)$ time by Lemmas 28, 29, and 30, respectively. Then, a $(1|1)_R$-centroid can be computed as the local optimal point of L^* in $O(n \log^2 n)$ time by Lemma 14. Combining with the $O(n^2 \log n)$-time preprocessing for computing the angular sorted sequence $P(v)$s and the bounding box enclosing $\mathcal{C}(V)$, we have the following theorem.

Theorem 3 *The* $(1|1)_R$*-centroid problem can be solved in* $O(n^2 \log n)$ *time.*

5 Concluding Remarks

In this paper, we revisited the (1|1)-centroid problem on the Euclidean plane under the consideration of minimal distance constraint between facilities, and proposed an $O(n^2 \log n)$-time algorithm, which close the bound gap between this problem and its unconstrained version. Starting from a critical observation on the medianoid solutions, we developed a pruning tool with indefinite region remained after pruning, and made use of it via multi-level structured parametric search approach, which is quite different to the previous approach in [5,11].

Considering distance constraint between facilities in various competitive facility location models is both of theoretical interest and of practical importance. However, similar constraints are rarely seen in the literature. It would be good starting points by introducing the constraint to the facilities between players in the $(r|X_p)$-medianoid and $(r|p)$-centroid problems, maybe even to the facilities between the same player.

References

1. R. Cole, "Slowing down sorting networks to obtain faster sorting algorithms," *Journal of the ACM*, vol. 34, no. 1, pp. 200–208, 1987.
2. R. Cole, "Parallel merge sort," *SIAM Journal on Computing*, vol. 17, no. 4, pp. 770–785, 1988.
3. A. Dasci, "Conditional Location Problems on Networks and in the Plane," In: H. A. Eiselt, V. Marianov (eds.), *Foundations of Location Analysis*, Springer, New York, pp. 179–206, 2011.
4. I. Davydov, Y. Kochetov, and A. Plyasunov, "On the complexity of the $(r|p)$-centroid problem in the plane," *TOP*, vol. 22, no. 2, pp. 614–623, 2013.
5. Z. Drezner, "Competitive location strategies for two facilities," *Regional Science and Urban Economics*, vol 12, no. 4, pp. 485–493, 1982.
6. Z. Drezner, and Z. Eitan, "Competitive location in the plane," *Annals of Operations Research*, vol. 40, no. 1, pp. 173–193, 1992.
7. H. A. Eiselt and G. Laporte, "Sequential location problems," *European Journal of Operational Research*, vol. 96, pp. 217–231, 1997.
8. H. A. Eiselt, G. Laporte, and J.-F. Thisse, "Competitive location models: a framework and bibliography," *Transportation Science*, vol. 27, pp. 44–54, 1993.
9. H. A. Eiselt, V. Marianov, Vladimir and T. Drezner, Tammy, "Competitive Location Models," In: G. Laporte, S. Nickel, F. Saldanha da Gama (eds.), *Location Science*, Springer International Publishing, pp. 365–398, 2015.
10. S. L. Hakimi, "On locating new facilities in a competitive environment," *European Journal of Operational Research*, vol. 12, no. 1, pp. 29–35, 1983.
11. S. L. Hakimi, "Locations with spatial interactions: competitive locations and games," In: P.B. Mirchandani, R.L. Francis (eds), *Discrete location theory*, Wiley, New York, pp. 439–478, 1990.
12. P. Hansen, J.-F. Thisse, and R. W. Wendell, "Equilibrium analysis for voting and competitive location problems," In: P. B. Mirchandani, R. L. Francis RL (eds), *Discrete location theory*, Wiley, New York, pp 479–501, 1990.

13. H. Hotelling, "Stability in competition," *Economic Journal*, vol. 39, 41–57, 1929.
14. D. T. Lee, Y. F. Wu, "Geometric complexity of some location problems," *Algorithmica*, vol. 1, no. 1, pp. 193–211, 1986.
15. N. Megiddo, "Applying parallel computation algorithms in the design of serial algorithms," *Journal of the ACM*, vol. 30, no. 4, pp. 852–865, 1983.
16. N. Megiddo, "Linear-time algorithms for linear programming in R^3 and related problems," *SIAM Journal on Computing*, vol. 12, no. 4, pp. 759–776, 1983.
17. F. Plastria, "Static competitive facility location: an overview of optimisation approaches." *European Journal of Operational Research*, vol. 129, no. 3, pp. 461–470, 2001.
18. A. Reiser, "A linear selection algorithm for sets of elements with weights," *Information Processing Letters*, vol. 7, no. 3, pp. 159–162, 1978.
19. D. R. Santos-Peñate, R. Suárez-Vega, and P. Dorta-González, "The leader–follower location model," *Networks and Spatial Economics*, vol. 7, no. 1, pp. 45–61, 2007.